

REMARKS/ARGUMENTS

Claims 1-25 are pending in this application, all of which stand rejected as a result of the March 17, 2003 Office Action. Specifically, claims 1, 2, 7, 9, 10, 21, 22, 23, and 25 have been rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,279,074 (Pence); claims 11, 12, 14, and 15 have been rejected under section 102(e) as being anticipated by U.S. Patent No. 6,272,605 (Le); claims 3, 4, 5, 6, 8, and 24 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Pence in view of U.S. Patent No. 6,490,666 (Cabrera); claim 13 has been rejected under section 103(a) as being unpatentable over Le in view of Pence; and claims 16-20 have been rejected under section 103(a) as being unpatentable over Le in view of Pence and Cabrera. Following entry of the amendment, claim 23 has been amended.

Applicants respectfully disagree with, and traverse, the stated grounds for rejection. As described below, the claims recite various features that are neither taught nor suggested by the prior art. Applicants thus submit that all of the claims are novel and non-obvious over the prior art of record.

Below, applicants discuss various claim features that define over the prior art of record:

Using a separate queue for each medium (claims 1, 12, 19, 21, and 25)

Claim 1 is directed to a method of recalling data objects stored on a plurality of media, and calls for:

- "creating a plurality of queues, wherein each one of said queues corresponds to one of said media" and

- "placing said requests on the created queues, wherein each request is placed on the queue corresponding to the medium on which the requested data object is located".

In other words, claim 1 calls for a separate queue for each medium, and requires that recall requests are placed on a particular queue depending on which medium the requested data object is located. None of the prior art cited by the Examiner teaches this feature.

Pence teaches that a single queue stores requests for objects located on different tapes.

As explained by Pence:

FIG. 4 illustrates a queue of recall requests When the host system 2 receives or generates another recall request, that new recall request is placed in the recall queue 50. The new recall request may be directed toward a volume on a storage device currently mounted in a drive or in a storage device not mounted. ... If the host system 2 receives a recall request toward a storage device 11 that is not mounted, that request too is placed in the recall queue 50.

[Pence, col. 6, ll. 1-14.] In other words, Pence clearly states that there is one queue (50) that contains all of the recall requests generated by a particular host, regardless of where the requested data is located. Claim 1, however, does not call for a single queue to store all recall requests, but instead requires a separate queue for each medium. Since this claim feature is not identically taught by Pence, Pence does not anticipate claim 1.

Additionally, claims 12, 19, 21, and 25 recite similar features. In particular:

- Claim 12 calls for first and second queues corresponding to first and second media;
- Claim 19 calls for a plurality of queues, each corresponding to a different medium, and inserting a request into a particular queue based on which medium the requested item is located;
- Claim 21 calls for a plurality of queues corresponding to a plurality of media,

and queuing each of a plurality of requests on a queue corresponding to the medium on which the requested data object is located; and

- Claim 25 calls for placing requests on a plurality of queues depending on which medium the requested data is located on.

As described below, none of these features is taught by any of the prior art cited by the Examiner.

Claims 21 and 25 have been rejected as being anticipated by Pence. Both of these claims recite features similar to those called for by claim 1. Thus, for the reasons discussed above in connection with claim 1, Pence does not anticipate claims 21 and 25.

Claim 12 has been rejected as being anticipated by Le. Le does not teach the use of first and second queues corresponding to first and second media. Le teaches the use of a single "recall queue" (col. 6, ll. 52-56). The Examiner has read the claimed "first queue" onto Le's "recall queue," but then reads the second queue onto "a second host recall request" (see Office Action, p. 10). However, a "second host recall request" is not the same thing as a "second queue." Evidently, the Examiner has assumed that if Le places recall requests from a first host on the recall queue, then Le would place recall requests from a second host onto a different queue. Le, however, contains no support for this assumption. Le teaches the use of (one) recall queue, and there is absolutely no teaching or suggestion in Le of a second queue. Thus, the Examiner's position on claim 12 represents an unwarranted and unsupported extension of Le.

Claim 19 has been rejected over Le in view of Pence. Both the present application and Cabrera are assigned to Microsoft Corporation, and both have an inventor (Pudipeddi) in

common. Thus, under 35 U.S.C. § 103(c), Cabrera cannot properly be used to form the basis for an obviousness rejection.

Pence and Le do not— either alone or in combination — teach or suggest the feature of using a separate queue for each one of a plurality of media, or the feature of placing a recall request on a particular queue based on which medium the requested object is located on. Moreover, Cabrera cannot be used to establish obviousness against this application assigned to Microsoft Corporation. Since claims 1, 12, 19, 21, and 25 all recite the features discussed above, these claims are neither anticipated by, nor obvious over, the prior art cited by the Examiner.

Two sequences in a queue (claims 2, 16, and 24)

Claim 2 calls for: “organizing the requests on each queue in a first and a second sequence, each sequence comprising a set of requests whose offsets are monotonically increasing within the respective sequence.” Similarly, claim 16 calls for “first and second sequences ... wherein the offsets of the requests within said first sequence are a maximally monotonically increasing series,” and claim 24 calls for “first and second sequences, the offsets of the data objects requested in each of said first and second sequences comprising a maximally monotonically increasing series.” These features (which applicants shall refer to as the “two sequences” feature) are not taught or suggested by the prior art of record.

The “two sequences” feature is illustrated, by way of example, in FIGS. 8A through 8E of the present application. Because many storage media are sequential access media (e.g., tapes), recall requests in a queue may be organized based on the order in which the requested data objects appear on the medium — i.e., in order of those objects’ offsets into the medium.

Thus, the reading head can move forward from one request to the next without having to shuttle back and forth through the tape. However, since requests may be added to the queue while the tape is being read, it is possible that a request will be added to the queue for a data object that is located behind the current position of the reading head. Thus, two sequences are maintained: one for items ahead of the current position of the reading head, and one for items located behind the current position. Thus, the tape drive moves forward through the tape to process the first sequence, rewinds, and then moves forward through the tape to process the second sequence. Since these sequences are monotonically increasing, it is not necessary to move the tape backward in the middle of a processing a sequence. (See application, p. 21, line 6 through p. 24, line 20).

Pence and Le do not teach the use of a queue having first and second sequences with respect to the order of the offsets of requested items, as called for by claims 2, 16, and 24 (and, as noted above, Cabrera cannot be used to establish obviousness of the present claims under 35 U.S.C. § 103(c)). In rejecting claims 2, 16, and 24, the Examiner has read the “two sequences” feature onto Pence, col. 2, l. 64 through col. 3, l. 2, which states:

The recall operation for the next recall request to the first storage device is performed after determining that the next recall request has the higher priority. The first storage device is demounted from the drive after determining that the recall request to the second storage device has the higher priority.

This passage does not teach the “two sequences” feature. Pence organizes recall requests in order of priority, and may call for tapes to be mounted, dismounted, rewound, etc., in order to process recall requests in the specified priority. Claims 2, 16, and 24, on the other hand, call for requests to be organized in two sequences based on the order in which the requested objects appear on a medium. Thus, the “two sequences” feature places requests in an order

based on where the requested objects appear on a medium, while the cited passage of Pence orders requests according to a priority and *without regard* to the order in which the objects appear on the medium. Thus, the cited passage of Pence does not anticipate the “two sequences” feature of claims 2, 16, and 24.

Additionally, the Examiner has read the “two sequences” feature of claim 24 onto col. 10, ll. 37-41 of Pence. This additional cited portion of Pence is not a teaching, but rather a fragment of a claim, which states:

a queue data structure including a plurality of recall requests to the first storage device and priority data associated with each recall request in the queue data structure to recall data from one of the first and second storage devices mounted in the drive.

It is unclear why the Examiner believes this claim fragment teaches first and second sequences. The quoted claim fragment mentions nothing about sequences within a queue, or the ordering of recall requests in a sequence based on the offset at which the requested item is located. While the claim fragment mentions “first and second storage devices,” these are clearly not “first and second sequences” as recited in claim 24.

In addition to the fact that the cited portion of Pence fails to teach the “two sequences” feature, applicants note that the “two sequences” feature is also not taught by Le and Pence. Moreover, applicants note that the Examiner has not proposed any interpretation of, or modification to, Le and Pence that would yield that “two sequences” feature.

As demonstrated above, the “two sequences” feature is neither anticipated by, nor obvious over, the prior art cited by the Examiner, and the rejection of claims 2, 16, and 24 should be withdrawn.¹

¹ Additionally, it should be noted that claim 18 calls for a new request to be inserted in such a manner that the monotonically increasing nature of the second sequence is preserved. For the reasons discussed herein,

Concurrent recall of two data items (claim 11)

Claim 11 calls for two data objects to be identified and concurrently recalled. The Examiner asserts that claim 11 is anticipated by Le. The portions of Le cited by the Examiner have nothing to do with concurrent recall of two data objects.

The Examiner has cited col. 9, lines 44-48 of Le with regard to the claimed feature of concurrently using first and second drives to retrieve first and second data objects from first and second media. In particular, the cited portion of Le states:

The logic of FIG. 6 allows a second intervening host to recall over a lower priority first requesting host and, at the same time, allow the first requesting host to perform recalls should the second intervening host fail or be slow to respond.

Evidently, the Examiner has misinterpreted the phrase “at the same time” in the above-quoted passage. It is clear from context that the above-quoted passage has nothing to do with allowing two hosts to perform recalls at the same time. Rather, the above-quoted passage (as well as FIG. 6 in general) addresses the problem of how to resolve conflicting priorities of recall requests. The phrase “at the same time” refers to the fact that the method of FIG. 6 takes into account two competing concerns – i.e., the method of FIG. 6 allows a second host to have its high priority requests processed ahead of a first hosts low priority requests, but will allow the first host to perform lower priority recalls if the second host is slow to respond. The above quoted passage, however, does not in any way suggest that these high and low priority recalls could proceed concurrently, as called for in claim 11.

Thus, the reading of claim 11 onto the cited portions of Le is erroneous, and the section 102 rejection of claim 11 based on that reading should be withdrawn.

Summary of the 102 and 103 Rejections

As demonstrated above, the Examiner has not demonstrated that claims 1, 2, 11, 12, 21, and 25 are anticipated by the prior art, and has not demonstrated that claims that claims 16, 19, and 24 are obvious over the prior art. According, applicants request that the rejection of these claims be withdrawn. Moreover, since claims 3-10, 13-20, and 22-24 depend, either directly or indirectly from at least one of the above-mentioned claims, these claims are patentable at least by reason of their dependency.

Thus, applicants submit that the section 102 and 103 rejections have been overcome as to all claims, and request that these rejections be withdrawn.

Amendment to Claim 23

Claim 23 has been amended to correct a typographical oversight. The amendment to claim 23 does not affect the scope of the claim, and is not made for a reason related to patentability. No new matter is introduced.

Drawings

Formal drawings were filed with this application. The Examiner has not acknowledged acceptance of the formal drawings. It is requested that, in the next Office

DOCKET NO.: MSFT-0232/160298.1
Application No.: 09/783,820
Office Action Dated: March 17, 2003

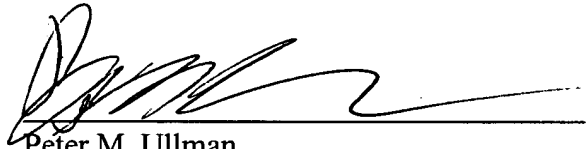
PATENT

Action, the Examiner acknowledge that the formal drawings filed with the application are acceptable.

Conclusion

For all of the foregoing reasons, applicants respectfully request that the Examiner reconsider and withdraw the rejection of claims 1-25. Applicants submit that this case is in condition for allowance, and request an early Notice of Allowance.

Respectfully submitted,



Peter M. Ullman
Registration No. 43,963

Date: July 16, 2003

Woodcock Washburn LLP
One Liberty Place - 46th Floor
Philadelphia PA 19103
Telephone: (215) 568-3100
Facsimile: (215) 568-3439